arXiv:1604.06770v2 [cs.DB] 25 Jul 2016

A Hybrid Approach to Query Answering under Expressive Datalog$^{\pm}$

Mostafa Milani[1], Andrea Calì[2,3], and Leopoldo Bertossi[1]

[1]School of Computer Science
Carleton University, Canada

[2]Dept of Computer Science
Birkbeck, Univ. of London, UK

[3]Oxford-Man Institute of Quantitative Finance
University of Oxford, UK

{mmilani, bertossi}@scs.carleton.ca
andrea@dcs.bbk.ac.uk

Abstract. Datalog$^{\pm}$ is a family of ontology languages that combine good computational properties with high expressive power. Datalog$^{\pm}$ languages are provably able to capture many relevant Semantic Web languages. In this paper we consider the class of weakly-sticky (WS) Datalog$^{\pm}$ programs, which allow for certain useful forms of joins in rule bodies as well as extending the well-known class of weakly-acyclic TGDs. So far, only nondeterministic algorithms were known for answering queries on WS Datalog$^{\pm}$ programs. We present novel deterministic query answering algorithms under WS Datalog$^{\pm}$. In particular, we propose: *(1)* a bottom-up grounding algorithm based on a query-driven chase, and *(2)* a hybrid approach based on transforming a WS program into a so-called sticky one, for which query rewriting techniques are known. We discuss how our algorithms can be optimized and effectively applied for query answering in real-world scenarios.

1 Introduction

The Datalog$^{\pm}$ family of ontology languages [4], which extends Datalog with explicit existential quantification, has been gaining importance in the area of ontology-based data access (OBDA) due to its capability of capturing several conceptual data models and Semantic Web languages as well as offering efficient query answering services in many variants relevant for applications.

The core feature of Datalog$^{\pm}$ languages are the so-called existential rules (a.k.a. *tuple-generating dependencies* or *TGDs*). Such rules allow the inference (entailment) of new atoms from an initial set of ground atoms, typically a database, through the *chase* procedure [7,9,11]. For example, consider the rule $\forall X \forall Y\ r(X,Y) \rightarrow \exists Z\, s(X,Z)$ (in the following we shall omit universal quantifiers, keeping only the existential ones) and a database D constituted by a single atom $r(a,b)$. A chase step will generate (notice that X and Y correspond to a and b respectively in this step) the atom $s(a,\zeta_1)$, where ζ is a *labelled null*, that is, a placeholder for an unknown value; notice that the constant b is lost

in this step as it doesn't appear in the new atom. In general the chase may not terminate. Answering a conjunctive query (CQ) q under a database D and a set of TGDs Σ amounts to computing all atoms entailed by $D \cup \Sigma \cup q$ (where q is seen as a TGD) or, equivalently, evaluating q over all atoms entailed by $D \cup \Sigma$.

Consider the following simple example. Let $\Sigma = \{\sigma_1, \sigma_2\}$; σ_1 is the rule $emp(X) \rightarrow \exists Y\ rep(X, Y)$, asserting that each employee reports to someone; σ_2 is the rule $rep(X, Y) \rightarrow mgr(Y)$, asserting that anyone to whom someone else reports is a manager. Now assume $D = \{mgr(ann), emp(joe)\}$. Let us consider the conjunctive query q_1 defined as $q_1(W_1) \leftarrow rep(W_1, W_2)$, where q_1 is now a new predicate (we use the head predicate q_1 for convenience, but any predicate name would suit). This query, which is itself a TGD[1], which we denote with δ, asks for all those who report to someone. Clearly the plain evaluation $q_1(D)$ over D returns no answer ($q_1(D) = \varnothing$), as D says nothing about who reports to whom. However, we are to reason on atoms entailed by $D \cup \Sigma \cup \delta$; therefore, we apply the chase of D w.r.t. $\Sigma \cup \delta$. By applying σ_1 to $emp(joe)$ we obtain $rep(joe, \zeta_1)$, where ζ_1 is a labeled null (in logical terms this means $\exists Y\ rep(joe, Y)$); then by applying σ_2 to $rep(joe, \zeta_1)$ we obtain $mgr(\zeta_1)$; no more atoms can be entailed by the chase. Now, if we evaluate q_1 on all entailed atoms, thus computing all atoms entailed by $D \cup \Sigma \cup \delta$, we get the answer $\langle joe \rangle$; this is because $D \cup \Sigma \cup \delta$ entails that *joe* reports to someone: $D \cup \Sigma \cup \delta \models q_1(joe)$. Consider now the query q_2 defined as $q_2(W_1) \leftarrow mgr(W_1)$, asking for all those who are managers. In this case, evaluating q_2 over all atoms entailed by $D \cup \Sigma$ returns the answer $\langle ann \rangle$; notice that $D \cup \Sigma$ does not entail the answer $\langle \zeta_1 \rangle$ for q_2; this because ζ_1 is just a placeholder for an unknown value — we know that the manager to whom *joe* reports exists, but we do not know who he (or she) is.

Conjunctive query answering under general TGDs is undecidable; languages of the Datalog$^{\pm}$ family impose therefore restrictions on the form of rules so as to guarantee decidability and certain computational problems, most prominently, conjunctive query answering.

Guarded and *weakly-guarded* Datalog$^{\pm}$ (the latter generalising the former) were the first decidable Datalog$^{\pm}$ languages, inspired by guarded logic and characterised by the presence of a *guard atom* in each rule that contains all variables of that rule [5]. The *sticky* Datalog$^{\pm}$ [5] language was introduced to capture a "proper" notion of *join* in rules, that is, the occurrence of variables in two distinct atoms of a rule body in the absence of a guard for that rule. *Weakly-sticky (WS)* Datalog$^{\pm}$ extends sticky Datalog$^{\pm}$ by also capturing the well-known class of *weakly-acyclic* rules [9]. As an example, consider the following set of rules Σ, where some body variables are *marked* (by a hat sign, e.g. $\hat{X}$; see [5]; notice that X and $\hat{X}$ are the same variable) as the result of a procedure that identifies occurrences of variables (the marked ones) corresponding, in the chase procedure, to a value that can eventually be lost in some subsequent chase step.

$$\sigma_1 :\ v(X) \rightarrow \exists Y\ r(X, Y). \qquad \sigma_3 :\ r(X, \hat{Y}), r(\hat{Y}, Z) \rightarrow p(X, Z).$$
$$\sigma_2 :\ p(\hat{X}, \hat{Y}) \rightarrow \exists Z\ p(Y, Z). \qquad \sigma_4 :\ p(\hat{X}, Y), p(Y, Z) \rightarrow t(Y, Z).$$

[1] Here we adopt the usual notation for conjuntive queries, where the head appears on the left-hand side.

A set of TGDs is *sticky* if, for each rule, there is no marked variable in that rule that appears more than once in the body of the same rule. Intuitively, stickiness can be defined by means of the following (semantic) property: during a chase step according to a rule σ, each value corresponding to a variable appearing more than once in σ is not lost in the chase step, and it is also never lost in any subsequent step involving atoms where it appears. Notice that the set of rules Σ above is not sticky, as easily seen.

Weak acyclicity is defined using the notion of *rank* of a position, i.e. a predicate attribute. In a position of finite (resp., infinite) rank, the number of labelled nulls that can appear in the chase procedure is finite (resp., infinite). In the set of TGDs Σ above, $\Pi_F(\Sigma) = \{v[1], r[1], r[2]\}$ contains the positions with *finite rank* and $\Pi_\infty(\Sigma) = \{p[1], p[2], t[1], t[2]\}$ those with *infinite rank*. A Datalog$^\pm$ program is *weakly-acyclic* if all positions have finite rank. The set Σ of TGDs above is not weakly-acyclic.

A set of TGDs is WS if, for each TGD, every marked variable that appears more than once in the body also appears at least once in a finite-rank position. This notion generalizes both stickiness and acyclicity, because the stickiness condition applies to variables that appear *only* in positions with infinite rank. Notice that Σ above is WS. Specifically, in σ_3, the repeated variable Y appears in positions $r[1]$ and $r[2]$, which are in $\Pi_F(\Sigma)$. In σ_4, Y is repeated not marked.

To answer a conjunctive query q under a set Σ of TGDs (or other types of ontological rules) and a database D, two main approaches were proposed in the literature: *grounding* (or *expansion*) and *query rewriting*. In the grounding approach, variables are suitably replaced by constants (or nulls) in the body of a rule, so that the head of the rule yields a (ground) atom that is entailed by the program. The aforementioned chase is in fact a grounding procedure. The grounding allows the computation of all atoms entailed by $D \cup \Sigma$, onto which q can be then evaluated. In the rewriting approach, the query q is rewritten, according to Σ, into another query q_R (possibly in another language different from that of q), so that the correct answers can be obtained by evaluating q_R directly on D.

The rewriting approach is usually considered more efficient than the grounding because in the former only the query is manipulated, according to the rules, while the data is left unchanged; on the contrary, the grounding approach requires the expansion of the given data, whose size is normally much larger than that of the query and of the rules.

CQ answering can be done in polynomial time in data for WS programs. However, so far, no non-trivial deterministic algorithm for CQ answering has been devised for WS Datalog$^\pm$. In this paper we devise algorithms for the efficient implementation of conjunctive query answering under WS Datalog$^\pm$. Our contributions are as follows.

1. We propose a bottom-up technique, based on grounding, which is a variant of the chase procedure, and relies on a terminating chase-like procedure that is *resumed* a number of times that depends on the query to be evaluated. Once the procedure is resumed a sufficient number of times for the query,

the same query is evaluated together with the result of the procedure, that is a set of (ground) rules, yielding the correct answers to the query under the given WS set of TGDs.
2. We propose a *hybrid* approach between grounding and rewriting as follows. First, with a novel algorithm, we transform the given set Σ of TGDs into another, all whose positions in Π_F become of rank 0 (that is, positions where only constants can appear in the chase or grounding). Next, certain variables in such positions are *grounded*, that is, they are replaced by selected constants of the given instance. The obtained program is a sticky program, for which a well-known query rewriting technique for CQ answering can be applied. The rewriting (a union of CQs) is finally evaluated on the initial instance.

Both techniques we propose yield algorithms for CQ answering that are of the same data complexity as the lower-bound complexity of CQ answering under weakly-stick programs. Moreover, both our algorithms, unlike the one for WS Datalog$^{\pm}$ in [5], are deterministic. The advantage of the first, pure-grounding technique is that we can pre-compute off-line a ground Datalog program (possibly containing nulls, treated as constants), which serves for answering every query up to a certain number of variables by simply evaluating it on the minimum model of the above Datalog program (in fact, a relational instance). The second, hybrid algorithm relies on a partial grounding, again computed off-line, as well as on an on-the-fly rewriting of every query; the advantages of this approach are that *(a)* the grounding is query-independent and generally much smaller than a complete grounding; *(b)* the final step consists of a mere evaluation of a union of CQs on the given database. Notice that in both our approaches the last step can be performed by executing an SQL query, then offering the possibility of taking advantage of optimizations of RDBMSs. Several optimization strategies are possible for these algorithms (this is ongoing work). We argue that our techniques set the basis for efficient CQ answering under expressive Datalog$^{\pm}$ languages such as WS Datalog$^{\pm}$. A full version of this paper is available at [16].

2 Preliminaries

In this section, we review some basic notions which we use in the paper.

2.1 Basic definitions

We assume an infinite universe of data constants Γ_C, an infinite set of (labeled) nulls Γ_N, and an infinite set of variables Γ_V (used in rules and queries). We denote by uppercase letters (e.g. X, Y, Z) variables, while $\boldsymbol{X}$ is a sequence of variables $X_1, \ldots, X_k$ with $k \geq 0$. We use the same notation for sets of variables. A relational schema $\mathcal{R}$ is a finite set of relation names (or predicates). A position $p[i]$ identifies the i-th argument of a predicate p.

A homomorphism is a structure-preserving mapping $h : \Gamma_C \cup \Gamma_N \cup \Gamma_V \rightarrow \Gamma_C \cup \Gamma_N \cup \Gamma_V$ such that $c \in \Gamma_C$ implies $h(c) = c$. For atoms and conjunctions of atoms, we denote by Π-homomorphism a homomorphism that is the identity on the terms that appear in a set Π of positions.

A *tuple-generating dependency* (or *TGD*, also called *existential rule*) σ on a schema $\mathcal{R}$ is a formula $p_1(\boldsymbol{X}_1), \ldots, p_n(\boldsymbol{X}_n) \rightarrow \exists \boldsymbol{Y}\ p(\boldsymbol{X}, \boldsymbol{Y})$ in which $p, p_1, \ldots, p_n$ are predicates in $\mathcal{R}$ and $\boldsymbol{X} \subseteq \bigcup \boldsymbol{X}_i$ and $\bigcup \boldsymbol{X}_i$ are universal variables that are implicitly quantified. We denote by $\mathit{head}(\sigma)$ and $\mathit{body}(\sigma)$ the head atom $p(\boldsymbol{X}, \boldsymbol{Y})$ and the set of the body atoms $p_1(\boldsymbol{X}_1), \ldots, p_n(\boldsymbol{X}_n)$ of σ, respectively. Variables in $\boldsymbol{Y}$ are called *existential variables*. A rule is *ground* if its terms are in $\Gamma_C \cup \Gamma_N$.

An instance D for $\mathcal{R}$ is a (possibly infinite) set of atoms with predicates in $\mathcal{R}$ and arguments from $\Gamma_C \cup \Gamma_N$. A database D is an instance that contains only atoms with arguments from Γ_C. The *active domain* of a database D denoted by $\mathit{active}(D)$ is the set of constants that appear in D.

A rule σ is satisfied by an instance I, written $I \models \sigma$, if the following holds: whenever there exists a homomorphism h such that $h(\mathit{body}(\sigma)) \subseteq I$, then there exists a homomorphism h' as an extension of h that maps existential variables of σ into terms in $\Gamma_C \cup \Gamma_N$, such that $h'(\mathit{head}(\sigma)) \subseteq I$. An instance I satisfies a set Σ of TGDs, denoted $I \models \Sigma$, if $I \models \sigma$ for each $\sigma \in \Sigma$.

A conjunctive query (CQ) has the form $q(\boldsymbol{X}) \leftarrow p_1(\boldsymbol{X}_1), ..., p_n(\boldsymbol{X}_n)$ where $p_1, ..., p_n$ are predicate names in $\mathcal{R}$, q is a predicate name not in $\mathcal{R}$, $\boldsymbol{X} \subseteq \bigcup \boldsymbol{X}_i$ and the $\boldsymbol{X}_i$ are sequences of variables or constants. A Boolean CQ (BCQ) over $\mathcal{R}$ is a CQ having head predicate q of arity 0 (i.e., no variables in $\boldsymbol{X}$). The answer to a CQ $q(\boldsymbol{X}) \leftarrow p_1(\boldsymbol{X}_1), \ldots, p_n(\boldsymbol{X}_n)$ over an instance I, denoted as $q(I)$, is the set of all n-tuples $t \in \Gamma_C^n$ for which there exists a homomorphism h such that $h(p_1(\boldsymbol{X}_1), \ldots, p_n(\boldsymbol{X}_n)) \subseteq I$ and $h(\boldsymbol{X}) = t$. A BCQ has only the empty tuple $\langle\rangle$ as possible answer, in which case we say it has a positive answer, denoted $I \models q$.

A program $\mathcal{P}$ consists of a set of rules $\Sigma^{\mathcal{P}}$ and a database $D^{\mathcal{P}}$ over same schema $\mathcal{R}$. $\mathcal{P}$ is a ground program if $\Sigma^{\mathcal{P}}$ is a set of ground rules. Given a program $\mathcal{P}$ and a CQ q, the answers to q are those that are true in all models of $\mathcal{P}$. Formally, the models of $\mathcal{P}$, denoted as $\mathit{mods}(\mathcal{P})$, is the set of all instances I such that $I \supseteq D^{\mathcal{P}}$ and $I \models \Sigma^{\mathcal{P}}$. The answers to a CQ q over $\mathcal{P}$, denoted as $\mathit{ans}(q, \mathcal{P})$, is the set of n-tuples $\{t \mid t \in q(I), \forall I \in \mathit{mods}(\mathcal{P})\}$. The answer to a BCQ q is positive, denoted as $\mathcal{P} \models q$, if $\mathit{ans}(q, \mathcal{P})$ is not empty.

2.2 Chase and Grounding

The *chase* procedure is a fundamental algorithm in various database problems including implication of database dependencies, query containment and CQ answering under dependencies [3,11,14]. The chase has been broadly employed in CQ answering in the presence of dependencies [5,9]; the intuition is that, given a set of dependencies over a database schema and a fixed database instance as input, the chase "repairs" the instance so that the result satisfies the constraints. The result of the chase procedure, also called chase, is a so-called *universal model* [9], i.e., a representative of all models in $\mathit{mods}(\mathcal{P})$; therefore, the answers to a CQ q under dependencies (in the open-world assumption, also called *certain answers*), can be computed by evaluating q over the chase (and discarding the answers containing labeled nulls). The chase under TGDs, which we do not describe in detail, is in fact a form of *grounding*; in Section 3 we propose a grounding technique for answering CQs under WS Datalog$^{\pm}$ based on a variant of the chase.

2.3 Datalog$^{\pm}$ and the Stickiness Paradigm

Query answering with respect to a set of TGDs is generally undecidable as proved in [2]. The Datalog$^{\pm}$ family contains syntactic classes of TGDs that impose restrictions on the form of the rules to guarantee decidability and in many cases tractability of query answering. Two relevant decidability paradigms are *guardedness* and *stickiness*. In this paper we concentrate on stickiness, which is a syntactic condition on the join variables in the body of the rules.

Sticky programs Sticky rules are defined by means of a body variable *marking procedure* that takes as input the set $\Sigma^{\mathcal{P}}$ of rules. It has two steps:

1. *Preliminary step*: for each $\sigma \in \Sigma^{\mathcal{P}}$ and for each variable $X \in \mathit{body}(\sigma)$, if there is an atom $a \in \mathit{head}(\sigma)$ such that X does not appear in a, mark each occurrence of X in $\mathit{body}(\sigma)$.
2. *Propagation step*: for each $\sigma \in \Sigma^{\mathcal{P}}$, if a marked variable in $\mathit{body}(\sigma)$ appears at position π, then for every $\sigma' \in \Sigma^{\mathcal{P}}$ (including σ), mark each occurrence of the variables in $\mathit{body}(\sigma')$ that appear in $\mathit{head}(\sigma')$ in the same position π.

Example 1. Consider a program $\mathcal{P}$, with $\Sigma^{\mathcal{P}}$ as following set of rules in which the marked variables (denoted by hat signs) after applying the preliminary step:

$$r(X,Y), p(X,Z) \rightarrow s(X,Y,Z). \qquad u(X) \rightarrow \exists\, Y\ r(Y,X).$$
$$s(\hat{X},Y,\hat{Z}) \rightarrow u(Y).$$

In the first rule, variables X and Z are marked after applying one propagation step since they appear in the head in marked positions $(s[1], s[3])$, and the final, marked rules are:

$$r(\hat{X},Y), p(\hat{X},\hat{Z}) \rightarrow s(X,Y,Z). \qquad u(X) \rightarrow \exists\, Y\ r(Y,X).$$
$$s(\hat{X},Y,\hat{Z}) \rightarrow u(Y). \qquad \blacksquare$$

$\mathcal{P}$ is sticky when, at the end of the marking procedure over $\Sigma^{\mathcal{P}}$, there is no rule with a marked variable in its body that occurs more than once. From Example 1, we can see that the program is *not* sticky since X in the first rule is marked and occurs twice in $r[1]$ and $p[1]$.

Sticky programs enjoy *first-order rewritability*. A class of programs is first-order rewritable if, for every program $\mathcal{P}$ in the class, and for every BCQ q, there is a first-order query $q_{\mathcal{P}}$ such that, $\mathcal{P} \models q$ if and only if $D^{\mathcal{P}} \models q_{\mathcal{P}}$. Thus, under first-order rewritable programs, CQs can be answered by constructing the (finite) rewritten first-order query [10], and then evaluating it over the extensional database. Since evaluation of first-order queries is in AC_0 in data complexity [1], it immediately follows that CQ answering under first-order rewritable classes of rules, including sticky ones, is in AC_0.

Weakly-Sticky programs Weakly-sticky (WS) programs generalize sticky programs and the well known class of weakly-acyclic programs. The class of WS programs is not first-order rewritable but CQ answering over these programs is proved to be tractable in data complexity [5]. The definition of WS programs appeals to conditions on repeated variables in its rule bodies, and is based on the notion of *dependency graph* and the positions with finite rank in such a graph [9] that we explain in Example 2.

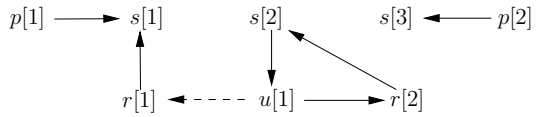

Fig. 1. Dependency graph for $\Sigma^{\mathcal{P}}$ in Example 2

Example 2. (Example 1 cont.) Consider the set of rules $\Sigma^{\mathcal{P}}$ in program $\mathcal{P}$. The dependency graph of $\Sigma^{\mathcal{P}}$ is a directed graph constructed as follows: The vertices in V are positions of the predicates in schema of $\Sigma^{\mathcal{P}}$, and the edges in E are defined as it follows. For every $\sigma \in \Sigma^{\mathcal{P}}$ and non-existential variable x in $head(\sigma)$ and in position π in $body(\sigma)$: (1) for each occurrence of x in position π' in $head(\sigma)$, create an edge from π to π'; (2) for each existential variable z in position π'' in $head(\sigma)$, create a *special edge* from π to π''. The dependency graph of $\Sigma^{\mathcal{P}}$ from Example 1 is illustrated in Figure 1. The special edge from $u[1]$ to $r[1]$ is shown by a dotted arrow that depicts values invention by existential variable Y in the last rule. The *rank* of a position is the maximum number of special edges over all (finite or infinite) paths ending at that position. Accordingly, $\Pi_F(\Sigma^{\mathcal{P}})$ denotes the set of positions of finite rank, and $\Pi_\infty(\Sigma^{\mathcal{P}})$ the set of positions of infinite rank. Intuitively, $\Pi_F(\Sigma^{\mathcal{P}})$ captures positions where finitely many values may appear during the chase; and $\Pi_\infty(\Sigma^{\mathcal{P}})$ those where infinitely many fresh null values may occur during the chase. In this example, $\Pi_\infty(\Sigma^{\mathcal{P}})$ is empty. The rank of $u[1], s[2], r[2], s[3], p[1]$ and $p[2]$ is zero and the rank of $r[1]$ and $s[1]$ is one. ■

A program $\mathcal{P}$ is WS if for every rule in $\Sigma^{\mathcal{P}}$ and every variable in its body that occurs more that once, the variable is either non-marked or appears at least once in a position in $\Pi_F(\Sigma^{\mathcal{P}})$ (position with finite rank).

Example 3. (Example 2 cont.) $\mathcal{P}$ is WS. $\Pi_F(\Sigma^{\mathcal{P}})$ contains $s[3]$, $p[1]$ and $p[2]$ and the other positions of the predicates in $\Sigma^{\mathcal{P}}$ are in $\Pi_\infty(\Sigma^{\mathcal{P}})$. The repeated variable X in the body of the second rule is marked, but it appears at least once in the finite-rank position $p[1]$. ■

3 Grounding Based on a Query-Driven Chase

In this section we adopt the chase procedure proposed in [15] for query answering under WS Datalog$^{\pm}$ programs as a basis for a query-driven grounding algorithm, called GroundWS. GroundWS takes as input a WS Datalog$^{\pm}$ program $\mathcal{P}$ and a CQ q and returns a ground program $\mathcal{P}'$ for which q can be efficiently answered.

The GroundWS algorithm uses the notion of applicability that we explain next. A rule σ and homomorphism h are *applicable* if: (a) $h(body(\sigma)) \subseteq \mathcal{H} \cup D^{\mathcal{P}}$ in which $\mathcal{H}$ is the set of head atoms of $\Sigma^{\mathcal{P}'}$ (the set of already grounded rules); (b) There is no atom $a \in \mathcal{H}$ and homomorphism h' as an extension of h such that h' maps existential variables in σ to fresh nulls in Γ_N and $h'(head(\sigma))$ is $\Pi_F(\Sigma^{\mathcal{P}})$-homomorphic to a. The second applicability condition is imposed to prevent infinite grounding steps as we describe next. Importantly, we use $\Pi_F(\Sigma^{\mathcal{P}})$-homomorphism instead of ordinary homomorphism to consider only the null values that appear in the positions with infinite rank, which is sufficient

to prevent infinite grounding steps. We now illustrate the algorithm GroundWS, which consists of the following steps:

1. Initialize $\Sigma^{\mathcal{P}'}$ to $\varnothing$ and $D^{\mathcal{P}'}$ to $D^{\mathcal{P}}$.
2. For every rule $\sigma \in \Sigma^{\mathcal{P}}$ applicable with homomorphism h, add $h'(\sigma)$ into $\Sigma^{\mathcal{P}'}$, in which h' is extension of h that maps existential variables of σ into fresh null values in Γ_N.
3. Apply Step 2 iteratively, until there are no more applicable rules.
4. Resume Step 2 after freezing every labeled null value in $\Sigma^{\mathcal{P}'}$, where by freezing a null we mean replacing it with a special constant in Γ_C, which henceforth is considered as a constant but is never returned in the result of a query. Repeat this resumption M_q times, where M_q is the number of variables in q.
5. Return $\mathcal{P}'$.

Notice that every pair of rule and homomorphism in Step 2 is applied only once. Moreover, if there are more than one pairs of applicable rule\homomorphism, then GroundWS applies them in a level saturating fashion. More specifically, GroundWS chooses the rule and the homomorphism for which the body atoms have the smallest maximum level. Here the level of an atom is 0 if it is in $D^{\mathcal{P}}$, and it is the maximum level of the body atoms of a rule plus one for the head atom of the rule. It is important to notice that by freezing a null value we consider it as a constant only for deciding homomorphic atoms (specifically in the second applicability condition), and not during query answering. That is, the frozen nulls still can not appear in query answers since they are not in the active domain of the extensional database.

Example 4. Consider WS program $\mathcal{P}$ with $D^{\mathcal{P}} = \{p(a,b), c(b)\}$, and a BCQ $q \leftarrow u(X)$, and a set of rules $\Sigma^{\mathcal{P}}$ as follows:

$$\sigma_1 \colon \quad p(X,Y) \;\rightarrow\; \exists Z\, p(Y,Z). \qquad\qquad \sigma_2 \colon \quad p(X,Y), c(X), p(Y,Z) \;\rightarrow\; u(Y).$$

We start from $D^{\mathcal{P}}$ and iteratively generate ground rules by mapping via homomorphism the body of the rules in $\Sigma^{\mathcal{P}}$ into $D^{\mathcal{P}}$ or the head of the rules in $\Sigma^{\mathcal{P}'}$ (the current set of ground rules). The basic algorithm is as follows: we iteratively add a ground rule to $\Sigma^{\mathcal{P}'}$ if its head atom is not homomorphic to the head of a rule already in $\Sigma^{\mathcal{P}'}$, until no new rule can be added. This "cautious" procedure, similar to the chase procedure in [13,15], guarantees termination. In our example, GroundWS stops after adding only one ground rule $\sigma_1' : p(a,b) \;\rightarrow\; p(b,\zeta_1)$ to $\Sigma^{\mathcal{P}'}$, where ζ_1 is a labelled null in Γ_N. In order to complete our grounding, we *resume* the above basic algorithm M_q times, where M_q is the number of variables in q. Before each resumption, we *freeze* the labelled nulls, which after having frozen are considered as constants. In our example, we have only another resumption, which adds the rules $\sigma_1'' : p(b,\zeta_1) \;\rightarrow\; p(\zeta_1,\zeta_2)$ and $\sigma_2' : p(b,\zeta_1), c(b), p(\zeta_1,\zeta_2) \;\rightarrow\; u(\zeta_1)$. Resumptions are needed, intuitively, to capture applications of rules in the chase procedure where a join variable, appearing in two or more distinct atoms, is mapped to a labelled null. ■

Notice that the number of resumptions depends on the query, which makes our grounding also dependent of the query; however, for practical purposes, we could ground with N resumptions so as to be able to answer queries with up to N existential variables, and if a query with more than N existential variables is to be answered, we can incrementally retake the already-computed grounding and add the required number of resumptions.

Theorem 1. For every WS program $\mathcal{P}$ and CQ q, GroundWS runs in PTIME with respect to the size of $D^{\mathcal{P}}$ and in 2EXPTIME with respect to the size of $\mathcal{P}$ and q, and returns a ground program $\mathcal{P}'$ such that, $ans(q, \mathcal{P}) = ans(q, \mathcal{P}')$.

Proof (sketch): GroundWS runs in polynomial time essentially because of the second applicability condition, which at each resumption prevents the generation in in $\Sigma^{\mathcal{P}'}$ of two distinct (ground) rules having $\Pi_F(\Sigma)$-homomorphic heads; since the number of terms in the positions of $\Pi_F(\Sigma)$ is polynomial with respect to $D^{\mathcal{P}}$ [9, Theorem 3.9], the PTIME membership follows. Notice that the above condition, within a resumption "phase", prevents the generation of some (ground) rules in $\Sigma^{\mathcal{P}'}$ that are necessary for answering q; these rules depend on replacing a join variable with a null value in the rule body; the subsequent resumption phase adds at least one of such rules to $\Sigma^{\mathcal{P}'}$. The 2EXPTIME combined complexity is because exponentially many terms can appear in the positions of $\Pi_F(\Sigma)$ with respect to the size of $\mathcal{P}$ and q. This is implicit in the proof of [9, Theorem 3.9]. As a result, the total number of atoms in each resumption phase is double exponential in the size of $\mathcal{P}$ and q.

The weak-stickiness of $\Sigma^{\mathcal{P}}$ implies that such null values continue to appear in the head of consequent ground rules all the way to the head atoms mapped to the query. Hence there are at most M_q such rules generated for answering q, given that each rule "saturates" one of the M_q existential variables in q. Therefore only M_q resumptions are necessary. □

4 Partial Grounding for Weakly-Sticky Datalog$^{\pm}$

In this section we propose a partial grounding algorithm, called PartialGroundingWS, that takes a WS Datalog$^{\pm}$ program $\mathcal{P}$ and transforms it into a sticky Datalog$^{\pm}$ program $\mathcal{P}'$ such that $\mathcal{P}'$ is equivalent to $\mathcal{P}$ for CQ answering. PartialGroundingWS selectively replaces certain variables in positions of finite rank with constants from the active domain of the underlying database. Our algorithm requires that the set of rules $\Sigma^{\mathcal{P}}$ in the input program satisfies the condition that there is no existential variable in $\Sigma^{\mathcal{P}}$ in any finite-rank position; therefore each position in the input set of rules $\Sigma^{\mathcal{P}}$ will have rank either 0 or ∞. The reason for this requirement is the convenience of grounding variables at zero-rank positions by replacing them by constants rather than by labeled nulls. This does not really restrict the input programs since, as we will show later, an arbitrary program can be transformed by the ReduceRank algorithm to a program that has the requirement.

Before illustrating the PartialGroundingWS algorithm, we therefore present an algorithm, called ReduceRank, that takes a program $\mathcal{P}$ and compiles it into an equivalent program $\mathcal{P}'$ whose rule set $\Sigma^{\mathcal{P}'}$ has only zero-rank or infinite-rank positions. The ReduceRank algorithm is inspired by the reduction method in [12] for transforming a weakly-acyclic program into an existential-free Datalog program. Given a program $\mathcal{P}$, ReduceRank executes of the following steps.

1. Initialize $\Sigma^{\mathcal{P}'}$ to $\Sigma^{\mathcal{P}}$ and $D^{\mathcal{P}'}$ to $D^{\mathcal{P}}$.
2. Choose a rule σ in $\Sigma^{\mathcal{P}'}$ with an existential variable in a position with rank 1. Notice that if there are existential variables in positions with finite rank, at least one of the positions has rank 1.
3. Skolemize σ as σ' by replacing the existential variable with a functional term. For example, $\sigma : p(X, Y) \rightarrow \exists Z\ r(Y, Z)$, becomes $\sigma' : p(X, Y) \rightarrow r(X, f(X))$.
4. Replace the Skolemized predicate r that has the function term with a new expanded predicate of higher arity (the arity of r plus 1) and introduce a *fresh* special constant of Γ_C to represent the function symbol. The new constant precedes its arguments in a newly introduced position. For example, $r(X, f(X))$ becomes, $r'(X, \mathsf{f}, X)$; we are therefore expanding the position $r[2]$.
5. Replace the expanded predicate in other rules and analogously expand other predicates in positions where variables appearing in the expanded position appear: if a variable appears in $\sigma \in \Sigma^{\mathcal{P}}$ in an expanded position π and also in another position π_1, then also π_1 is expanded (with its predicate) and the same variables of π also appear in π_1, thus preserving the join. More precisely, let X be a variable at position π in an atom $r(\ldots, X, \ldots)$ that is expanded in some rule σ into $r(\ldots, X_1, X_2, \ldots)$; if there is another rule $\sigma_1 \in \Sigma^{\mathcal{P}}$ containing atoms $r(\ldots, Y, \ldots)$ and $s(\ldots, Y, \ldots)$ such that Y appears in the former at position π_1 and in the latter at position π_2, then the first atom is replaced by $r'(\ldots, Y_1, Y_2, \ldots)$ (expanded version on π_1 with the new variable Y) and the second one is replaced by $s'(\ldots, Y_1, Y_2, \ldots)$ (expanded version on π_2). For example, $r(X, Y), t(Y, Z) \rightarrow s(X, Y, Z)$ becomes $r'(X, Y, Y'), t'(Y, Y', Z) \rightarrow s'(X, Y, Y', Z)$. Notice that if a predicate is expanded in a head-atom in a position where an existential variable occurs, the new positions are not required and are filled with the special symbol $\triangle$.
6. If the expanded predicates have extensional data, add new rules to $\Sigma^{\mathcal{P}'}$ to "load" the extensional data into the expanded predicates. For example, if r has extensional data, we add a rule, $r(X, Y) \rightarrow r'(X, Y, \triangle)$. Here, $\triangle$ is used to fill the new position in the expanded predicate since it does not carry extensional data.
7. Repeat Steps 2 to 6 until there is no existential variable in a finite-rank position.

Note that in Step 3 only the body variables that also appear in the head participate as arguments of the function term. For example, in the Skolemization of $p(X, Y) \rightarrow \exists Z\ r(Y, Z)$, the function term does not include X since the rule can be broken down into $p(X, Y) \rightarrow u(Y)$ and $u(Y) \rightarrow \exists Z\ r(Y, Z)$. Notice that

given a CQ q over $\mathcal{P}$, Steps 2 to 6 are also applied on q obtaining a new CQ q' over $\mathcal{P}'$.

Example 5. Let $\mathcal{P}$ be a program with $\Sigma^{\mathcal{P}}$ as follows.

$$\sigma_1 : \ v(X) \ \rightarrow \ \exists Y \ r(X,Y). \qquad \sigma_3 : \ t(X,Y), v(X) \ \rightarrow \ p(X,Y).$$
$$\sigma_2 : \ r(X,Y) \ \rightarrow \ \exists Z \ t(X,Z). \qquad \sigma_4 : \ p(X,Y) \ \rightarrow \ \exists Z \ p(Y,Z).$$

In this program, $\Pi_F(\Sigma^{\mathcal{P}}) = \{v[1], r[1], r[2], t[1], t[2]\}$. ReduceRank will eliminate Y in σ_1 and Z in σ_2, but not Z in σ_4 since the later is in an infinite rank position. ReduceRank chooses Y in σ_1 over Z in σ_2 since, Y is in $r[2]$ with zero rank and Z is in $t[2]$ with rank 1. After applying Steps 2-6, σ_1 and σ_2 become $\sigma_1' : \ v(X) \ \rightarrow \ r'(X, \mathsf{f}, X)$ and $\sigma_2' : \ r'(X,Y,Y') \ \rightarrow \ \exists Z \ t(X,Z)$. By removing Y from σ_1, Z in σ_2 is placed in a position with zero rank. ReduceRank repeats Steps 2-6 to eliminate Z in σ_2 which results into $\Sigma^{\mathcal{P}'}$:

$$v(X) \rightarrow r'(X, \mathsf{f}, X). \qquad r'(X,Y,Y') \rightarrow t'(X, \mathsf{g}, X).$$
$$t'(X,Y,Y'), v(X) \rightarrow p'(X, \triangle, Y, Y'). \qquad p'(X,X',Y,Y') \rightarrow \exists Z \ p'(Y,Y',Z, \triangle).$$

Notice that ReduceRank does not try to remove Z in the last rule, since it is in the infinite rank position $p[3]$. Note also that p is expanded twice since both its positions can host labeled nulls generated by Z in σ_2. ■

Proposition 1. Given a CQ q over a program $\mathcal{P}$, ReduceRank runs in EXPTIME with respect to the size of $\Sigma^{\mathcal{P}}$ and returns a CQ q' over a program $\mathcal{P}'$ such that $\mathcal{P}'$ has no existential variable in finite rank positions of $\Sigma^{\mathcal{P}'}$ and $ans(q, \mathcal{P}) = ans(q', \mathcal{P}')$.

For every rule in $\Sigma^{\mathcal{P}}$, there is only one corresponding rule in $\Sigma^{\mathcal{P}'}$. There are also rules in $\Sigma^{\mathcal{P}'}$ for loading the extensional data of the expanded predicates. Therefore, the number of rules in $\Sigma^{\mathcal{P}}$ is the same order of the size of $\Sigma^{\mathcal{P}'}$. The arity of the predicates in $\Sigma^{\mathcal{P}'}$ can have an exponential increase with respect to the arity of predicates in $\Sigma^{\mathcal{P}}$ which makes ReduceRank run in EXPTIME. $\mathcal{P}$ and $\mathcal{P}'$ are equivalent since the expanded predicate that represent the propagation of null values in $\mathcal{P}$ are applied on every possible rule in $\Sigma^{\mathcal{P}'}$ in Steps 5-6 .

Notice that ReduceRank preserves the weak-stickiness property because the property only concerns repeated marked variables that occur in infinite rank positions, while ReduceRank involves finite rank positions and it does not create a new marked variable or a new infinite rank position to break the property. We can therefore state the following.

Lemma 1. The class of WS programs is closed under ReduceRank.

Now that we explained the ReduceRank algorithm, we continue and present the PartialGroundingWS algorithm. Given a WS program $\mathcal{P}$, let us call *weak rules* the rules of $\Sigma^{\mathcal{P}}$ in which some repeated marked body variables (which we call weak variables) appear at least once in a position with finite rank. PartialGroundingWS transforms $\mathcal{P}$ into a sticky program $\mathcal{P}'$. The sticky program $\mathcal{P}'$ has the

same database as $\mathcal{P}$ ($D^{\mathcal{P}'} = D^{\mathcal{P}}$) and its set of rules $\Sigma^{\mathcal{P}'}$ is obtained by replacing the weak variables of $\Sigma^{\mathcal{P}}$ with every constants from the active domain of $D^{\mathcal{P}}$. Example 6 illustrates the PartialGroundingWS algorithm.

Example 6. Consider a WS program $\mathcal{P}$ with $D^{\mathcal{P}} = \{p(a,b), r(a,b)\}$ and $\Sigma^{\mathcal{P}}$ consisting of the following rules:

$$\sigma_1 : p(\hat{X}, \hat{Y}) \ \rightarrow \ \exists Z \ p(Y,Z). \qquad \sigma_3 : s(\hat{X}, Y, Z), r(\hat{X}, Y) \rightarrow t(Y,Z).$$

$$\sigma_2 : p(\hat{X}, Y), p(Y,Z) \ \rightarrow \ s(X,Y,Z).$$

Here σ_3 is a weak rule with X as its weak variable. Notice that Y in σ_2 and σ_3 are not weak since they are not marked (the hat signs show the marked variables). We replace X with constants a and b from $D^{\mathcal{P}}$. The result is a set of sticky rules $\Sigma^{\mathcal{P}'}$ that includes σ_1 and σ_2 as well as the following rules, $\sigma'_3 : \ s(a,Y,Z), r(a,Y) \rightarrow t(Y,Z)$ and $\sigma''_3 : \ s(b,Y,Z), r(b,Y) \rightarrow t(Y,Z)$. ■

Theorem 2. Let $\mathcal{P}$ be a WS program such that there is no existential variable in $\Sigma^{\mathcal{P}}$ in a finite rank position. PartialGroundingWS runs in polynomial time with respect to the size of $D^{\mathcal{P}}$ and it transforms $\mathcal{P}$ into a sticky program $\mathcal{P}'$ such that for every CQ q, the following holds: $ans(q, \mathcal{P}) = ans(q, \mathcal{P}')$.

Proof (sketch): $\mathcal{P}'$ is sticky since every weak variable, that by its definition breaks the weak-stickiness, is grounded. $\mathcal{P}$ and $\mathcal{P}'$ are equivalent for query answering since the weak variables of $\Sigma^{\mathcal{P}}$ are replaced in $\Sigma^{\mathcal{P}'}$ with every possible constant from $D^{\mathcal{P}}$, and based on our assumption on $\Sigma^{\mathcal{P}}$, only constants can substitute these variables. Additionally, PartialGroundingWS runs in polynomial time since weak variables are replaced with constants from $D^{\mathcal{P}}$. □

A possible optimization for PartialGroundingWS is to narrow down the values for replacing the weak variables, that is to ignore those constants in the active domain of $D^{\mathcal{P}}$ that can not appear in the positions where weak variables appear during the chase of $\mathcal{P}$. In Example 6, σ'_3 is not useful since a can never be assigned to X in σ_3. For this purpose, GroundWS can be applied to compute the possible values for partial grounding. For example, a CQ, $s(X,Y,Z), r(X,Y) \ \rightarrow \ q_g(X)$ returns constants for grounding the weak variable X in σ_3.

5 A Hybrid Approach

In this section we propose a query answering algorithm for WS programs based on a hybrid approach that combines ReduceRank and PartialGroundingWS from the previous section with a query rewriting algorithm for sticky programs [10]. Given a WS program $\mathcal{P}$ and a CQ q, hybrid query answering proceeds as follows:

1. Use ReduceRank to compile $\mathcal{P}$ into a WS program $\mathcal{P}'$ with no existential variable in finite rank positions.
2. Apply PartialGroundingWS on $\mathcal{P}'$ that results to a sticky program $\mathcal{P}''$.

3. Rewrite q into a first-order query q' using the rewriting algorithm proposed in [10] and answer q' over $D^{\mathcal{P}''}$ (any other sound and complete rewriting algorithm for sticky programs is also applicable at this step).

Example 7. Consider a WS program $\mathcal{P}$ with database $D = \{v(a)\}$ and $\Sigma^{\mathcal{P}}$ consisting of the following rules:

$$\begin{array}{llll} \sigma_1: & p(X,Y) \rightarrow \exists Z\, p(Y,Z). & \sigma_4: & r(X,Y), s(X,Z) \rightarrow c(Z). \\ \sigma_2: & p(X,Y), p(Y,Z) \rightarrow u(Y). & \sigma_5: & c(X) \rightarrow \exists Y\, p(X,Y). \end{array}$$
$$\sigma_3\ v(X) \rightarrow \exists Y\, r(X,Y).$$

The ReduceRank method removes the existential variable Y in σ_3. The result is a WS program $\mathcal{P}'$ with $\Sigma^{\mathcal{P}'}$:

$$\begin{array}{ll} p(X,Y) \rightarrow \exists Z\, p(Y,Z). & r'(X,Y,Y'), s(X,Z) \rightarrow c(Z). \\ p(X,Y), p(Y,Z) \rightarrow u(Y). & c(X) \rightarrow \exists Y\, p(X,Y). \end{array}$$
$$v(X) \rightarrow r'(X, \mathsf{f}, X).$$

Next, PartialGroundingWS grounds the only weak variable, X in σ'_4 with constant a which results into sticky program $\mathcal{P}''$ with $\Sigma^{\mathcal{P}''} = \{\sigma_1, \sigma_2, \sigma'_3, \sigma''_4, \sigma_5\}$, in which $\sigma''_4: r'(a,Y,Y'), s(a,Z) \rightarrow c(Z)$. $\mathcal{P}''$ is sticky and a CQs can be answered by rewriting it in terms of $\Sigma^{\mathcal{P}''}$ and answered directly on $D^{\mathcal{P}''} = D^{\mathcal{P}}$. ■

Corollary 1. Given a WS program $\mathcal{P}$ and a CQ q, the set of answers obtained from the hybrid approach is $ans(q, \mathcal{P})$.

6 Conclusions

WS Datalog$^{\pm}$ is an expressive ontology language with good computational properties and capable of capturing the most prominent Semantic Web languages. We proposed two deterministic algorithms for answering conjunctive queries on WS Datalog$^{\pm}$. In the first algorithm, a variant of the well-known chase, which proceeds in terminating "resumptions", generates an expansion of the given database that contains all (ground) atoms needed to answer the query; the expansion depends on the query as the number of resumptions is the number of existentially quantified variables of the query. For practical purposes, one can expand up to m resumptions off-line, and the expansion will serve to answer all queries with up to m existential variables. If at a certain point a query with more than m existential variables is to be processed, more resumptions can be performed from the expansion already computed. This, of course, if there are no changes in the given database. The second algorithm transforms a WS program into a sticky one by means of *(a)* a Skolemization and annotation procedure, which turns all finite-rank positions into zero-rank ones, followed by *(b)* a partial grounding on the zero-ranked positions. Then, the rewriting technique for WS programs is employed; the rewriting, which is in the language of union of conjunctive queries, is then evaluated directly on the given database.

Efficiency. Both algorithms we propose achieve the optimal lower bound in data complexity (i.e., in complexity calculated having only the database as input) for CQ answering under WS Datalog$^{\pm}$, that is PTIME. In the first algorithm, the expansion is computed *off-line*, and the final query processing step is a simple evaluation of a CQ on an instance. In the second algorithm, the rewriting is intensional (i.e., it does involve the data) and the final step is the evaluation of a union of CQs on the given database, which can easily done, for example, by evaluating an SQL query on the database.

In the light of the above considerations, we believe that our contribution sets the basis for practical query answering algorithms in real-world scenarios. We plan to continue our work by running experiments on large data sets. We also intend to refine the hybrid algorithm by limiting the number of CQs in the final rewriting; to do so, we will avoid the grounding of rules when we discover that having certain constants in certain position will not yeld any new atom; such discovery can be performed by analyzing the dependency graph and the TGDs in general. This refinement will improve the efficiency of the algorithm.

References

1. Abiteboul, S., Hull, R. and Vianu, V. *Foundations of Databases*. Addison-Wesley, 1995.
2. Beeri, C. and Vardi, M. Y. The implication problem for data dependencies. *Proc. Colloquium on Automata, Languages and Programming*, 1981, pp. 73–85.
3. Beeri, C. and Vardi, M. Y. A Proof Procedure for Data Dependencies. *Journal of ACM*, 1984, 31(4):718–741.
4. Calì, A., Gottlob, G. and Lukasiewicz, T. A General Datalog-based Framework for Tractable Query Answering over Ontologies. *Journal of Web Semantics*, 2012, 14:57–83.
5. Calì, A., Gottlob, G. and Pieris, A. Towards more Expressive Ontology Languages: The Query Answering Problem. *Artificial Intelligence*, 2012, 193:87–128.
6. Calì, A., Gottlob, G. and Kifer, M. Taming the Infinite Chase: Query Answering under Expressive Relational Constraints. *Journal of AI Research*, 2013, 48(1):115–174.
7. Deutsch, A., Nash, A. and Remmel, J. B. The Chase Revisited. *Proc. PODS*, 2008, pp. 149-158.
8. Duschka, O. M. and Genesereth, M. R. Answering Recursive Queries Using Views. *Proc. PODS*, 1997, pp. 109–116.
9. Fagin, R., Kolaitis, P. G., Miller, R. J. and Popa, L. Data Exchange: Semantics and Query Answering. *TCS*, 2005, 336:89–124.
10. Gottlob, G., Orsi, G. and Pieris, A. Query Rewriting and Optimization for Ontological Databases. *Proc. TODS*, 2014, 39(3):25.
11. Johnson, D. S. and Klug, A. Testing Containment of Conjunctive Queries under Functional and Inclusion Dependencies. *Proc. PODS*, 1984, pp. 164–169.
12. Krötzsch, M. and Rudolph, S. Extending Decidable Existential Rules by Joining Acyclicity and Guardedness. *Proc. IJCAI*, 2011, pp. 963–968.
13. Leone, N., Manna, M., Terracina, G. and Veltri, P. Efficiently Computable Datalog$^{\pm}$ Programs. *Proc. KR*, 2012, pp. 13-23.

14. Maier, D., Mendelzon, A. and Sagiv, Y. Testing Implications of Data Dependencies. *Proc. TODS*, 1979, pp. 152–152.
15. Milani, M. and Bertossi, L. Tractable Query Answering and Optimization for Extensions of Weakly-Sticky Datalog+-. *Proc. AMW*, 2015.
16. Milani, M., Calì, A. and Bertossi, L. A Hybrid Approach to Query Answering under Expressive Datalog$^{\pm}$. Tech. rep. available at `https://goo.gl/edg9FK`.

A Appendix

Proof of Theorem 1: GroundWS runs in PTIME because of the second condition (Condition (b)) for applicable rule-homomorphism that is used in Step 2 of the algorithm. Specifically, the condition prevents from Σ' multiple ground rules with $\Pi_F(\Sigma)$-homomorphic head atoms. Note that there are polynomially many values in the positions of $\Pi_F(\Sigma)$ during the algorithm. This follows from Theorem 3.9 in [9] that shows during the chase of a program $\mathcal{P}$, there are polynomially many values in the positions of $\Pi_F(\Sigma^{\mathcal{P}})$ with respect to the size of $D^{\mathcal{P}}$. As a result, the number of possible non-$\Pi_F(\Sigma^{\mathcal{P}})$-homomorphic atoms and also the number of possible ground rules at each resumption are polynomial with respect to the size of $D^{\mathcal{P}}$. Since the number of resumptions M_q is independent of the data, the total number of ground rules in GroundWS is also polynomial with respect to the size of $D^{\mathcal{P}}$ which proves the tractability of GroundWS.

The second condition for applicable rule-homomorphism may prevent some ground rules from Σ' that are necessary for answering q. Specifically, let σ and h be a rule-homomorphism that is not applicable because $b = h'(\mathit{head}(\sigma))$ (with h' extending h with mapping for $\exists$-variables in σ) is $\Pi_F(\Sigma^{\mathcal{P}})$-homomorphic to an atom $a = \mathit{head}(\sigma')$ with $\sigma' \in \Sigma'$. That means a and b differ in some null values, e.g. ζ_a and ζ_b resp that appear in non-$\Pi_F(\Sigma^{\mathcal{P}})$ positions. The null value ζ_b might appear in the head of the other ground rules in Σ' and by preventing $h'(\sigma)$ (with the head b that contains ζ_b) from Σ', we miss the query answers that depend (directly in the query's body or indirectly by the other applicable rules) on replacing a join variable with ζ_b.

In order to solve this issue, we use resumptions at Step 4. More specifically, a resumption freezes ζ_a and ζ_b mentioned above by replacing them with fresh constants everywhere in Σ'. Consequently, any join that depends on replacing a join variable with ζ_b is enabled by augmenting Σ' with $h'(\sigma)$. These joins might be in the body of the query or in the body of rules. In both cases the number of such joins is limited by M_q, the number of existential variables in the query. That is because of the weak-stickiness of $\mathcal{P}$ since the values that replace the join variables in non-$\Pi_F(\Sigma^{\mathcal{P}})$ positions continue to appear in the consequent head atoms and specifically those head atoms that are mapped to the body of the query. Therefore, after M_q resumptions Σ' contains every ground rule that is needed for answering q. ■

Proof of Proposition 1: To prove that ReduceRank runs in EXPTIME with respect to the size of $\Sigma^{\mathcal{P}}$ we show that each iteration of Steps 2-6 runs in EXPTIME. Let consider one such iteration that transforms $\mathcal{P}_i$ into $\mathcal{P}_{i+1}$ and

removes an existential variable Y_i in the rule σ_i and expands a predicate p_i in the head of σ_i to p'_i. Then, $\mathcal{P}_1$ is the input program $\mathcal{P}$ and $\mathcal{P}_n$ is the result program $\mathcal{P}'$ in which n is the total number of iterations that is equal to the total number of existential variables in positions with finite rank.

Y_i is in a position with the rank 1 and so expanding p_i does not expand any predicate in the body of σ_i during Step 5. As a result, every position in the program only gets expanded once during an iteration. Let r and b the number of rules and the maximum number of body atoms in $\Sigma^{\mathcal{P}}$ resp. that are not changed by the iterations. Also, let w_i be and the maximum arity of atoms in $\Sigma^{\mathcal{P}}$ at the i-th iteration. Then, the arity of p'_i (the expanded predicate of p_i) is at most increased by $b \times w_i$ (the maximum possible number of variables in the body of σ_i). Therefore, after propagating the expanded position in other rules, the maximum arity of the predicates $w_{i+1} = b \times w_i^2$ (each of the w_i positions is expanded once by $b \times w_i$). This means after n iterations the maximum arity of the predicates is $b^n \times w_i^{2n}$. On the other hand, n is at most $w_1 \times r$ and as the result $b^n \times w_i^{2n}$ is exponential with respect to w_1 and r and so the size of $\Sigma^{\mathcal{P}}$ which proves our claim.

In order to prove the correctness of the transformation, i.e. $ans(q, \mathcal{P}) = ans(q', \mathcal{P}')$, we prove the correctness of each iteration, $ans(q_i, \mathcal{P}_i) = ans(q_{i+1}, \mathcal{P}_{i+1})$. We do that by constructing a model $I_{i+1} \models \mathcal{P}_{i+1} \cup q_{i+1}$ for every model $I_i \models \mathcal{P}^i \cup q_i$. Let assume that removing Z_i introduces a function symbol f. We first build an assignment μ_i that maps a null value in I_i into a list of terms also in I_i. If σ_i is the rule containing Z_i, for every homomorphism h_i that maps both the body and the head of σ_i into I_i, μ maps $h_i(Z_i)$ into the list of terms that appear in $h_i(body(\sigma_i))$.

Now for every atom $a = p(t_1, ..., t_n) \in I_i$ we add an atom a' into I_{i+1} that is constructed as follows: (a) if p is not expanded in $\mathcal{P}_i$ then $a' = a$, (b) if p is expanded into p' in its k-th position, there are two possibilities, t_k is either a null value or a constant. If t_k is a constant expand it into $t_k, \triangle, ..., \triangle$ to fill the expanded positions, and if t_k is a null value, expand t_k into $\mu(t_k)$.

I_{i+1} is a model of $\mathcal{P}_{i+1} \cup q_{i+1}$ because for every homomorphism g that maps the body of a rule σ' into I_{i+1} we can make an extension of g' using μ that maps the head also into I_{i+1}.

In the other direction for any model I_{i+1} of $\mathcal{P}_{i+1} \cup q_{i+1}$, we build a model I_i of $\mathcal{P}_i \cup q_i$. This is simply by replacing any extended predicate p' with its original predicate p and removing additional terms, i.e. removing $\triangle$ symbols and the function symbols and their consequent terms with null values. For example, $p'(a, b, \mathsf{f}, c)$ becomes $p(a, b, \zeta)$ if p expanded by one, and $p'(a, b, d, \triangle)$ becomes $p(a, b, d)$. Again, it is straightforward to prove that I_i is a model of $\mathcal{P}_i \cup q_i$ by showing that for any homomorphism that maps the body of a rule into I_i there is an extension of it that maps the head into I_i. ■

Proof of Lemma 1: To prove the class of weakly-sticky programs is closed under ReuceRank, we show it is closed under each iteration of ReuceRank, i.e. Steps 2-6, that removes one existential variable in a finite rank position. Let $\mathcal{P}_i$

be the inputs of the iteration and $\mathcal{P}_{i+1}$ be the result of the iteration. Now, we prove two claims about $\mathcal{P}_{i+1}$:

1. Let π be a position in $\mathcal{P}_i$ with finite rank and π' the same position or one of the positions resulted by expanding π in $\mathcal{P}_{i+1}$. If π has finite rank, π' also has finite rank. Using proof by contradiction, assume π' is not finite, then there is a cycle in the dependency graph of $\mathcal{P}_{i+1}$ that includes π and passes a special edge. The same positions in the cycle or their corresponding positions in $\mathcal{P}_i$, if they are expanded, make a cycle in the dependency graph of $\mathcal{P}_i$ that includes π and also passes a special edge. Therefore, π is not finite which contradicts the assumption and proves the claim.
2. If a body variable X in $\mathcal{P}^i$ is not marked, the same variable (or the variables $X_1, ..., X_n$ if the position of X is expanded in $\mathcal{P}_{i+1}$) is not marked in $\mathcal{P}_{i+1}$. Considering the preliminary step of the marking procedure, if a body variable in $\mathcal{P}_i$ appears in the head, the corresponding body variable(s) in $\mathcal{P}_{i+1}$ also appear in the head of the corresponding rule. Similarly at the propagation step, if a body variable is not in marked positions in head of a rule in $\mathcal{P}_i$, the same variable or its corresponding variables in $\mathcal{P}_{i+1}$ do not appear in the head of the corresponding rule in $\mathcal{P}_{i+1}$ in a marked position. Therefore, if a variable is not marked during the steps of the marking procedure in $\mathcal{P}_i$, it is will not get marked in $\mathcal{P}_{i+1}$ either.

As a result of the claims, if there is no repeated marked variable in $\mathcal{P}_i$ that does not appear in finite rank positions, then there is no such variable in $\mathcal{P}_{i+1}$. Since each iteration preserves the weak-stickiness property, the whole transformation also preserves the property. ■

Proof of Theorem 2: $\mathcal{P}'$ is sticky since every weak variable that breaks the weak-stickiness property is replaced with constants. Also, $ans(q, \mathcal{P}) = ans(q', \mathcal{P}')$ follows from the correctness of Proposition 1 and the fact that the weak variables are replaced with every possible constant from D. It is important that no null value can appear in the positions of the weak variables in $\Sigma_{0,\infty}$ due to the first phase.

The algorithm runs in polynomial time with respect to the size of the database because running the first phase (ReduceRank) is data independent and the second phase (partial grounding) replaces variables with polynomially many values from the database. ■